The Effect of Block-wise Feedback on the Throughput-Delay Trade-off in Streaming

Gauri Joshi
EECS Dept., MIT
Cambridge, MA 02139, USA
Email: gauri@mit.edu

Yuval Kochman
School of CSE, HUJI
Jerusalem, Israel
Email: yuvalko@cs.huji.ac.il

Gregory W. Wornell
EECS Dept., MIT
Cambridge, MA 02139, USA
Email: gww@mit.edu

arXiv:1403.0259v1 [cs.IT] 2 Mar 2014

***Abstract*—Unlike traditional file transfer where only total delay matters, streaming applications impose delay constraints on each packet and require them to be *in order*. To achieve fast in-order packet decoding, we have to compromise on the throughput. We study this trade-off between throughput and in-order decoding delay, and in particular how it is affected by the frequency of block-wise feedback to the source. When there is immediate feedback, we can achieve the optimal throughput and delay simultaneously. But as the feedback delay increases, we have to compromise on at least one of these metrics. We present a spectrum of coding schemes that span different points on the throughput-delay trade-off. Depending upon the delay-sensitivity and bandwidth limitations of the application, one can choose an appropriate operating point on this trade-off.**

I. Introduction

A. Motivation

A recent report [1] shows that 62% of the Internet traffic in North America comes from real-time streaming applications such as NetFlix (28.88%) and YouTube (15.43%). Streaming traffic consumes such a large fraction of Internet bandwidth because video files inherently have a larger size than other forms of data. Thus, there is a need to develop transmission schemes which can ensure a high quality of experience to the user, with efficient use of available bandwidth.

Unlike traditional file transfer where only total delay matters, streaming imposes delay constraints on each individual packet. Further, many applications require in-order playback of packets at the receiver. Packets received out of order are buffered until the missing packets in the sequence are successfully decoded. In audio and video applications some packets can be dropped without affecting the streaming quality. However, other applications such as remote desktop, and collaborative tools such as Dropbox and Google Docs have strict order constraints on packets, where packets represent instructions that need to be executed in order at the receiver.

To ensure that packets are decoded in order, the transmission scheme must give higher priority to older packets that were delayed, or received in error due to channel noise. However, repeating old packets instead of transmitting new packets results in a loss in the overall rate at which packets are delivered to the user, that is, the throughput. Thus there is a fundamental trade-off between throughput and in-order decoding delay.

The throughput loss incurred to achieve low in-order decoding delay can be significantly reduced if the source receives feedback about packet losses, and thus can adapt its future transmission strategy to strike the right balance between old and new packets. We study this interplay between feedback and the throughput-delay trade-off.

B. Previous Work

Only a few papers in literature have analyzed streaming codes. Fountain codes [2] are capacity-achieving erasure codes, but they are not suitable for streaming because the decoding delay is proportional to the size of the data. Streaming codes without feedback for constrained channels such as adversarial and cyclic burst erasure channels were first proposed in [3], and also extensive explored in [4], [5]. The thesis [3] also proposed codes for more general erasure models and analyzed their decoding delay. Decoding delay has also been analyzed studied in [6], [7] in a multicast scenario with immediate feedback to the source.

However, decoding delay does not capture *in order* packet delivery which is required for streaming applications. This aspect is captured in the delay metrics in [8] and [9], which consider that packets are played in-order at the receiver. The authors in [8] analyze the throughput-delay trade-off for uncoded packet transmission over a channel with long feedback delay. In [9] we propose coding schemes that minimize playback delay in point-to-point streaming for the no feedback and immediate feedback cases. However, the case of block-wise feedback to the source remains to be explored.

C. Our Contributions

In this paper we consider this unexplored problem of how to effectively utilize block-wise feedback to the source to ensure in-order packet delivery to the user. In contrast to playback delay considered in [8] and [9], we propose a more versatile delay metric called the in-order decoding exponent. This metric captures the burstiness in the in-order decoding of packet for applications which require packets in-order, but do not necessarily play them at a constant rate.

When there is immediate feedback, we can achieve the best throughput-delay trade-off. But when the feedback comes in blocks, we have to compromise on the throughput to ensure fast in-order decoding. We present a spectrum of coding

schemes that span different points on the throughput-delay trade-off. Depending upon the delay-sensitivity, and bandwidth limitations of the application, one can choose an appropriate operating point on this trade-off. The proposed codes can be shown to be optimal over a broad class of schemes for the no feedback, and small feedback delay cases.

II. Problem Setup

A. System Model

We consider a point-to-point packet streaming scenario where the source has a large stream of packets $s_1, s_2, \cdots, s_n$. The encoder creates a coded packet $y_n = f(s_1,\ s_2\ ..s_n)$ in each slot n and transmits it over the channel. The encoding function f is known to the receiver. For example, if y_n is a linear combination of the source packets, the coefficients are included in the transmitted packet so that the receiver can use them to decode the source packets from the coded combination. Without loss of generality, we can assume that y_n is a linear combination of the source packets.

We consider an i.i.d. packet erasure channel where every transmitted packet is correctly received with probability p, and otherwise received in error and discarded. An erasure channel is a good model when encoded packets have a set of checksum bits that can be used to verify with high probability whether the received packet is error-free.

The receiver application requires the stream of packets to be *in order*. Packets received out of order are buffered until the missing packets in the sequence are decoded. Due to this in-order property, the transmitter can stop including s_k in coded packets when it knows that the receiver can decode s_k once all s_i for $i < k$ are decoded. We refer such packets as "seen" packets. The notion of "seen" is defined formally as follows.

Definition 1 (Seen Packets)**.** *A packet s_k is said to be "seen" by the transmitter when it knows that a coded combination that only includes s_k and packets s_i for $1 \leq i \leq k$ is received successfully.*

We consider that the source receives block-wise feedback about channel erasures after every d slots. Thus, before transmitting in slot $kd + 1$, for all integers $k \geq 1$, the source knows about the erasures in slots $(k-1)d + 1$ to kd. It can use this information to adapt its transmission strategy in slot $kd + 1$. Block-wise feedback can be used to model a half-duplex communication channel where after every d slots of packet transmission, the channel is reserved for the receiver to send feedback about the status of decoding. Note that the feedback can be used to estimate p, the probability of success of the erasure channel, when it is unknown to the source.

B. Throughput and Delay Metrics

We consider two metrics to measure the quality of streaming, the throughput τ and in-order decoding exponent λ. The throughput is the rate at which "innovative" coded packets are received. A coded packet is said to be "innovative" if it is linear independent with respect to the coded packets received until then. The bandwidth required is proportional to $1/\tau$. The throughput captures the overall rate at which packets go through the channel, irrespective of the order. The *in-order* decoding aspect is captured by a metric called the in-order decoding exponent λ which is defined as follows.

Definition 2 (In-order Decoding Exponent)**.** *Let T be the time between two successive instants of decoding one or more packets in-order. Then the in-order decoding exponent λ is*

$$\lambda \triangleq -\lim_{n\to\infty} \frac{\log \Pr(T > n)}{n}. \tag{1}$$

The relation (1) can also be stated as $\Pr(T > n) \doteq e^{-n\lambda}$ where $\doteq$ stands for asymptotic equality defined in [10, Page 63]. The in-order decoding exponent captures the burstiness in packet decoding. For example, if the streaming application plays one in-order packet in every slot, and if there are b packets in the receiver buffer, then the probability of an interruption in playback is proportional to $e^{-\lambda b}$.

In this paper we analyze how the trade-off between τ and λ is affected by the block-wise feedback delay d. We first consider the extreme cases of immediate feedback $(d = 1)$ and no feedback $(d = \infty)$ in Section III and Section IV respectively. This gives us insights into the analysis of the (τ, λ) trade-off for general d in Section V.

III. Immediate Feedback

In the immediate feedback $(d = 1)$ case, the source has complete knowledge of past erasures before transmitting each packet. We can show that a simple automatic-repeat-request (ARQ) scheme is optimal in both τ and λ. In this scheme, the source transmits the lowest index unseen packet, and repeats it until the packet successfully goes through the channel.

Since a new packet is received in every successful slot, the throughput $\tau = p$, the success probability of the erasure channel. The ARQ scheme is throughput-optimal because the throughput $\tau = p$ is equal to the information-theoretic capacity of the erasure channel [10]. Moreover, it also gives the optimal the in-order decoding exponent λ because one in-order packet is decoded in every successful slot. To find λ, first observe that the tail distribution of the time T, the interval between successive in-order decoding instants is,

$$\Pr(T > n) = (1-p)^n \tag{2}$$

Substituting this in Definition 2 we get the exponent $\lambda = -\log(1-p)$. Based on this analysis of the immediate feedback case, we can find limits on the range of achievable (τ, λ) for any feedback delay d as follows.

Lemma 1. *The throughput and delay metrics (τ, λ) achievable for any feedback delay d lie in the region $0 \leq \tau \leq \rho$, and $0 \leq \lambda \leq -\log(1-p)$.*

Proof: When feedback is received after blocks of $d > 1$ slots, the source has less knowledge about past erasures than in the immediate feedback ($d = 1$) case. Thus, the (τ, λ) trade-off when $d > 1$ is always worse than $(\tau, \lambda) = (p, -\log(1-p))$ the optimal trade-off for the immediate feedback ($d = 1$) case. ■

IV. No Feedback

Now we consider the other extreme case ($d = \infty$), where there is no feedback to the source about channel erasures. We propose a coding scheme and prove that it gives the best (τ, λ) trade-off among a class of codes called full-rank codes which are defined as follows.

Definition 3 (Full-rank Codes)**.** *In slot n we transmit a linear combination of all packets s_1 to $s_{V[n]}$, where the coefficients are chosen from a large enough field such that the coded combinations are independent with high probability. We refer to $V[n]$ as the transmit index in slot n.*

Conjecture 1. *Since the packets are required in-order at the receiver, we believe that given transmit index $V[n]$, there is no loss of generality in including all packets s_1 to $s_{V[n]}$.*

Hence we believe that there is no loss of generality in restricting our attention to full-rank codes.

Theorem 1. *The optimal throughput-delay trade-off among full-rank codes is $(\tau, \lambda) = (r, D(r||p))$ for all $0 \le r \le p$. It is achieved by the coding scheme with $V[n] = \lceil rn \rceil$ for all n.*

The term $D(r||p)$ is the binary information divergence function which is defined for $0 < p, r < 1$ as

$$D(r||p) = r \log \frac{r}{p} + (1-r) \log \frac{1-r}{1-p}. \tag{3}$$

Note that as $r \to 0$, $D(r||p)$ converges to $-\log(1-p)$, which is the optimal λ, as given by Lemma 1.

To prove Theorem 1, we first show that the scheme with transmit index $V[n] = \lceil rn \rceil$ in time slot n achieves the trade-off $(\tau, \lambda) = (r, D(r||p))$. Then we prove the converse by showing that no other full-rank scheme gives a better trade-off.

Proof of Achievability: Consider the scheme with transmit index $V[n] = \lceil rn \rceil$, where r represents the rate of adding new packets to the transmitted stream. The rate of adding packets is below the capacity of the erasure channel. Thus it is easy to see that the throughput $\tau = r$. Let $E[n]$ be the number of combinations, or equations received until time n. It follows the binomial distribution with parameter p. All packets $s_1 \cdots s_{V[n]}$ are decoded when $E[n] \ge V[n]$. Define the event $G_n = \{E[n] < V[n] \text{ for all } 1 \le j \le n\}$, that there is no packet decoding until slot n. The tail distribution of time T between successive in-order decoding instants is,

$$\begin{aligned}
\Pr(T > n) &= \sum_{k=0}^{\lceil nr \rceil - 1} \Pr(E[n] = k) \Pr(G_n | E[n] = k), \\
&= \sum_{k=0}^{\lceil nr \rceil - 1} \binom{n}{k} p^k (1-p)^{n-k} \Pr(G_n | E[n] = k), \\
&\doteq \sum_{k=0}^{\lceil nr \rceil - 1} \binom{n}{k} p^k (1-p)^{n-k}, && (4) \\
&\doteq \binom{n}{\lceil nr \rceil - 1} p^{\lceil nr \rceil - 1} (1-p)^{n - \lceil nr \rceil + 1}, && (5) \\
&\doteq e^{-nD(r||p)}, && (6)
\end{aligned}$$

where in (4), we remove the $\Pr(G_n | E[n] = k)$ when we take the asymptotic equality $\doteq$ because, by the Generalized Ballot theorem from [11], we can show that $\Pr(G_n | E[n] = k)$ is $\omega(1/n)$. Hence it is sub-exponential and does not affect the exponent of $\Pr(T > n)$. In (5), we only retain the $k = \lceil nr \rceil - 1$ term from the summation because for $r \le p$, that term asymptotically dominates other terms. Finally, we use the Stirlings approximation of the binomial coefficient $\binom{n}{k} \approx e^{nH(k/n)}$ to obtain (6).

Hence we have proved that the scheme with $V[n] = \lceil rn \rceil$ achieves the throughput-delay trade-off $(\tau, \lambda) = (r, D(r||p))$. ∎

Proof of Converse: First let us show that the transmit index $V[n]$ of the optimal full-rank scheme should be non-decreasing in n. Given a scheme which does not satisfy the non-decreasing property, we can permute the order of transmitting the coded packets such that $V[n]$ is non-decreasing in n. Changing the order of the transmitted packets will not affect the throughput τ. And it can in fact improve the in-order decoding exponent λ because decoding can occur sooner when the initial coded packets include fewer source packets.

In the proposed scheme with $V[n] = \lceil rn \rceil$, we add new packets to the transmitted stream at a constant rate r. But in general a full-rank scheme can vary the rate of adding packets. Suppose it uses rate r_i for n_i slots for all $1 \le i \le L$, such that $\sum_{i=0}^{L} n_i = n$ and $\sum_{i=1}^{L} n_i r_i = nr$. Then, the tail distribution of time T between successive in-order decoding instants is,

$$\begin{aligned}
\Pr(T > n) &= \sum_{k=0}^{\lceil \sum_{i=1}^{L} n_i r_i \rceil - 1} \Pr(E[n] = k) \Pr(G_n | E[n] = k), && (7) \\
&\doteq \sum_{k=0}^{\lceil nr \rceil - 1} \binom{n}{k} p^k (1-p)^{n-k}, && (8) \\
&\doteq e^{-nD(r||p)}. && (9)
\end{aligned}$$

Varying the rate of adding packets affects the term $\Pr(G_n | E[n] = k)$ in (7), but it is still $\omega(1/n)$ and we can eliminate it when we take the asymptotic equality in (8). As a result, the in-order delay exponent is same as that if we had a constant rate r of adding new packets to the transmitted stream. Hence we have proved that no other full-rank scheme can achieve a better (τ, λ) trade-off than $V[n] = \lceil nr \rceil$ for all n. ∎

Fig. 1 shows the (τ, λ) trade-off for the immediate feedback and no feedback cases, with success probability $p = 0.6$. The optimal trade-off with any feedback delay d lies in between these two extreme cases.

V. General Block-wise Feedback

In Section III and Section IV we considered the extreme cases of immediate feedback ($d = 1$) and no feedback ($d = \infty$) respectively. We now analyze the (τ, λ) trade-off with general block-wise feedback delay of d slots. We restrict our attention to a class of coding schemes called time-invariant schemes, which are defined as follows.

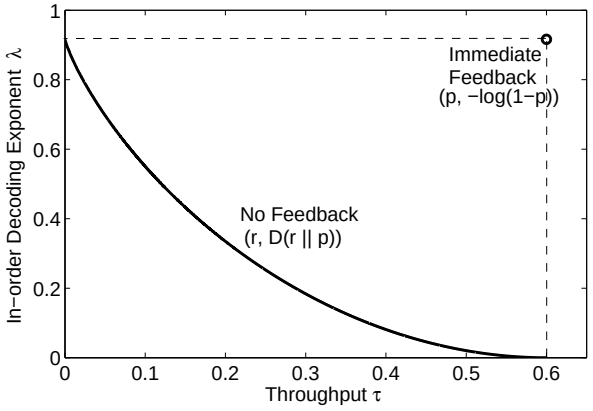

Fig. 1. The trade-off between in-order decoding exponent λ and throughput τ with success probability $p = 0.6$ for the immediate feedback $(d = 1)$ and no feedback $(d = \infty)$ cases.

Definition 4 (Time-invariant schemes)**.** *A time-invariant scheme is represented by a vector $\mathbf{x} = [x_1, \cdots x_d]$ where x_i, for $1 \leq i \leq d$, are non-negative integers such that $\sum_i x_i = d$. In each block we transmit x_i linear combinations of the i lowest-index unseen packets in the stream.*

The above class of schemes is referred to as time-invariant because the vector $\mathbf{x}$ is fixed across all blocks. Observe that as $d \to \infty$, the class of time-invariant schemes are equivalent to full-rank codes defined in Definition 3.

Conjecture 2. *For any coding scheme, there exists a corresponding time-sharing policy between time-invariant schemes that gives the same or strictly better (τ, λ) trade-off.*

We believe this conjecture is true because, it can be shown that any full-rank code can be expressed a time-sharing time-invariant scheme. By Conjecture 1 it follows that there is no loss of generality in focusing on time-invariant schemes.

There is also no loss of generality in restricting the length of the vector $\mathbf{x}$ to d. This is because we are still transmitting d independent coded packets. And adding fewer source packets to the coded combinations, can only increase the exponent λ.

A. Analyzing the (τ, λ) of time-invariant schemes

Given a vector $\mathbf{x}$, define p_d, as the probability of decoding the first unseen packet during the block, and S_d as the number of innovative coded packets that are received during that block. We can express $\tau_{\mathbf{x}}$ and $\lambda_{\mathbf{x}}$ in terms of p_d and S_d as,

$$(\tau_{\mathbf{x}}, \lambda_{\mathbf{x}}) = \left(\frac{\mathbb{E}[S_d]}{d}, -\frac{1}{d}\log(1-p_d) \right), \qquad (10)$$

where we get throughput $\tau_{\mathbf{x}}$ by normalizing the $\mathbb{E}[S_d]$ by the number of slots in the slots. We can show that the probability $\Pr(T > kd)$ of no in-order packet being decoded in k blocks is equal $(1-p_d)^k$. Substituting this in (1) we get $\lambda_{\mathbf{x}}$.

Example 1. *Consider the time-invariant scheme $\mathbf{x} = [1, 0, 3, 0]$ where block size $d = 4$. That is, we transmit 1 combination of the first unseen packet, and 3 combinations of the first 3 unseen packets. Fig. 2 illustrates this scheme for one channel realization. The probability p_d and $\mathbb{E}[S_d]$ are,*

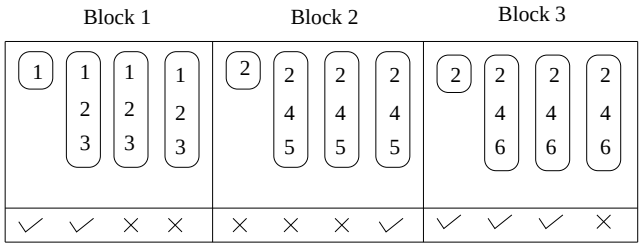

Fig. 2. Illustration of the time-invariant scheme $\mathbf{x} = [1, 0, 3, 0]$ with block size $d = 4$. Each bubble represents a coded combination, and the numbers inside it are the indices of the source packets included in that combination. The check and cross marks denote successful and erased slots respectively. The packets that are "seen" in each block are not included in the coded packets in future blocks.

$$p_d = p + (1-p)\binom{3}{3}p^3(1-p)^0 = p + (1-p)p^3, \quad (11)$$

$$\mathbb{E}[S_d] = \sum_{i=1}^{3} i \cdot \binom{4}{i} p^i (1-p)^{4-i} + 3p^4 = 4p - p^4, \qquad (12)$$

where in (12)*, we get i innovative packets if there are i successful slots for $1 \leq i \leq 3$. But if all 4 slots are successful we get only 3 innovative packets. We can substitute* (11) *and* (12) *in* (10) *to get the (τ, λ) trade-off.*

Remark 1. *Time-invariant schemes with different $\mathbf{x}$ can be equivalent in terms of the (τ, λ). In general, given $x_1 \geq 1$, if any $x_i = 0$, and $x_{i+1} = x \geq 1$, then the scheme is equivalent to setting $x_i = 1$ and $x_{i+1} = x - 1$, keeping all other elements of $\mathbf{x}$ the same. For example, $\mathbf{x} = [1, 1, 2, 0]$ gives the same (τ, λ) as $\mathbf{x} = [1, 0, 3, 0]$.*

B. Cost of Achieving Optimal τ or λ

In Section III we saw that for the immediate feedback case, we can achieve $(\tau, \lambda) = (p, -\log(1-p))$. However, when the feedback is delayed we can achieve optimal τ (or λ) only at the cost of sacrificing the optimality of the other metric. We now find the best achievable τ (or λ) with optimal λ (or τ).

Lemma 2 (Cost of Optimal Exponent λ)**.** *For a feedback delay of d slots, the best achievable throughput is $\tau = (1-(1-p)^d)/d$, when the in-order decoding exponent $\lambda = -\log(1-p)$.*

Proof: If we want to achieve $\lambda = -\log(1-p)$, we require p_d in (10) to be equal to $1-(1-p)^d$. The only scheme that can achieve this is $\mathbf{x} = [d, 0, \cdots, 0]$, where we transmit d copies of the first unseen packet. The number of innovative packets S_d received in every block is 1 with probability $1-(1-p)^d$, and zero otherwise. Hence, the best achievable throughput is $\tau = (1-(1-p)^d)/d$ with optimal $\lambda = -\log(1-p)$. ■

This result gives us insight on how much bandwidth (which is proportional to $1/\tau$) is needed for a highly delay-sensitive application which needs λ to be as large as possible.

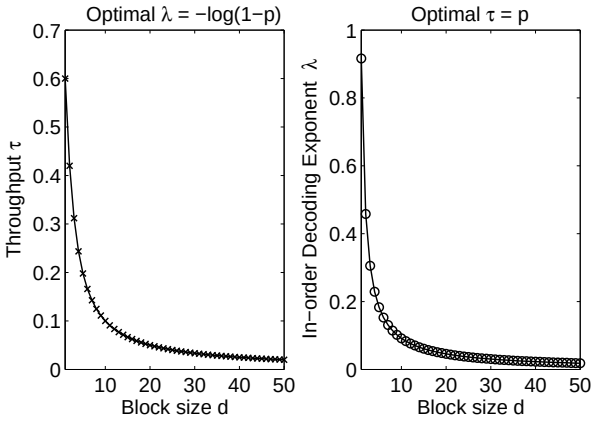

Fig. 3. Plot of the best achievable τ (or λ) versus d, while maintaining the optimal value of the other metric λ (or τ), for channel success probability $p = 0.6$.

Lemma 3 (Cost of Optimal Throughput τ)**.** *For a feedback delay of d slots, the best achievable in-order decoding exponent is $\lambda = -\log(1-p)/d$, when the throughput $\tau = p$.*

Proof: If we want to achieve $\tau = p$, we need to guarantee an innovation packet in every successful slot. The only time invariant scheme that achieve this is $\mathbf{x} = [1, 1, \cdots 1]$, and the vectors $\mathbf{x}$ that are equivalent to it as given by Remark 1. With $\mathbf{x} = [1, 1, \cdots 1]$, the probability of decoding the first unseen packet is $p_d = p$. Substituting this in (10) we get $\lambda = -\log(1-p)/d$, the best achievable λ when $\tau = p$. ∎

Fig. 3 shows the best achievable τ and λ versus d, when the other metric is at its optimal value. The plots in Fig. 3 correspond to moving leftwards and downwards respectively from the optimal trade-off $(p, -\log(1-p))$ in Fig. 1.

C. Finding Optimal (τ, λ) Trade-off

For any given throughput τ, our aim is to find the transmission scheme that achieves the maximum λ. We first prove that any convex combination of achievable points (τ, λ) can be achieved.

Theorem 2 (Convex Combinations of Time-invariant Schemes)**.** *Given time-invariant schemes $\mathbf{x}^{(i)}$ for $1 \leq i \leq B$, we can achieve the throughput-delay trade-off given by any convex combination of the points $(\lambda_{\mathbf{x}^{(i)}}, \tau_{\mathbf{x}^{(i)}})$ by time-sharing between the schemes.*

Proof: Here we prove the result for $B = 2$, that is time-sharing between two schemes. It can be extended to general B using induction. Given two time-invariant schemes $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ which achieve the throughput-delay trade-offs $(\lambda_{\mathbf{x}^{(1)}}, \tau_{\mathbf{x}^{(1)}})$ and $(\lambda_{\mathbf{x}^{(2)}}, \tau_{\mathbf{x}^{(2)}})$ respectively, consider a time-sharing strategy where, in each block we use the scheme $\mathbf{x}^{(1)}$ with probability μ and scheme $\mathbf{x}^{(2)}$ otherwise. Then, it is easy to see that the throughput on the new scheme is $\tau = \mu\tau_{\mathbf{x}^{(1)}} + (1-\mu)\tau_{\mathbf{x}^{(2)}}$.

Now we prove the in-order decoding exponent λ is also a convex combinations of $\lambda_{\mathbf{x}^{(1)}}$ and $\lambda_{\mathbf{x}^{(2)}}$. Let p_{d_1} and p_{d_2} be the probabilities of decoding the first unseen packet in a block using scheme $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ respectively. Suppose in an interval with k blocks, we use scheme $\mathbf{x}^{(1)}$ for h blocks, and scheme $\mathbf{x}^{(2)}$ in the remaining blocks, we have

$$\Pr(T > kd) = (1-p_{d_1})^h (1-p_{d_2})^{k-h}. \tag{13}$$

Using this we can evaluate λ as,

$$\lambda = \lambda_{\mathbf{x}^{(1)}} \lim_{k\to\infty} \frac{h}{k} + \lambda_{\mathbf{x}^{(2)}} \lim_{k\to\infty} \frac{k-h}{k} \tag{14}$$

$$= \mu\lambda_{\mathbf{x}^{(1)}} + (1-\mu)\lambda_{\mathbf{x}^{(2)}} \tag{15}$$

where we get (14) using (10). As $k \to \infty$, by the weak law of large numbers, the fraction h/k converges to μ. Hence, we have shown that we can interpolate between the (τ, λ) trade-off of two policies by time-sharing between them. ∎

The main implication of Theorem 2 is that, to find the optimal (τ, λ) trade-off, we only have to find the points $(\tau_{\mathbf{x}}, \lambda_{\mathbf{x}})$ that lie on the convex envelope of the achievable region spanned by all possible $\mathbf{x}$. We determine this optimal trade-off for $d = 2, 3$ in Lemma 4 and Lemma 5 below.

Lemma 4 (Optimal Trade-off for $d = 2$)**.** *The optimal (τ, λ) trade-off is the line joining points $\big((1-(1-p)^2)/2, -\log(1-p)\big)$ and $(p, -\log(1-p)/2)$.*

Proof: When $d = 2$ there are only two possible time-invariant schemes $\mathbf{x} = [2, 0]$ and $[1, 1]$ that give unique (τ, λ). By Remark 1, all other valid vectors $\mathbf{x}$ are equivalent to one of these schemes. From Lemma 2 and Lemma 3 we know that the (τ, λ) for these schemes are $((1-(1-p)^2)/2, -\log(1-p))$ and $(p, -\log(1-p)/2)$ respectively. By Theorem 2 we can achieve all (τ, λ) on the line joining these two points by time-sharing between the two policies. ∎

Lemma 5 (Optimal Trade-off for $d = 3$)**.** *The optimal (τ, λ) trade-off when $d = 3$ is the piecewise linear curve joining points*

$$(\tau_A, \lambda_A) = \left(\frac{1-(1-p)^3}{3}, -\log(1-p)\right), \tag{16}$$

$$(\tau_B, \lambda_B) = \left(\frac{2p(2-p)}{3}, -\frac{2}{3}\log(1-p)\right), \tag{17}$$

$$(\tau_C, \lambda_C) = \left(p, -\frac{\log(1-p)}{3}\right). \tag{18}$$

Proof: When $d = 3$ there are four time-invariant schemes $\mathbf{x}^{(1)} = [3, 0, 0], \mathbf{x}^{(2)} = [2, 1, 0], \mathbf{x}^{(3)} = [1, 2, 0]$ and $\mathbf{x}^{(4)} = [1, 1, 1]$ that give unique (τ, λ), as given by Definition 4 and Remark 1. From Lemma 2 and Lemma 3 we know that $(\tau_{\mathbf{x}^{(1)}}, \lambda_{\mathbf{x}^{(1)}}) = (\tau_A, \lambda_A)$ and $(\tau_{\mathbf{x}^{(4)}}, \lambda_{\mathbf{x}^{(4)}}) = (\tau_C, \lambda_C)$.

For the other two schemes, we first evaluate p_d and $\mathbb{E}[S_d]$ and substitute them in (10) to get, $(\tau_{\mathbf{x}^{(2)}}, \lambda_{\mathbf{x}^{(2)}}) = (\tau_B, \lambda_B)$ and

$$(\tau_{\mathbf{x}^{(3)}}, \lambda_{\mathbf{x}^{(3)}}) = \big((3p-p^3)/3, -(\log(1-p)^2(1+p))/3\big).$$

We can show that $\mathbf{x}^{(2)}$ gives a better trade-off than $\mathbf{x}^{(3)}$ by showing that for all p, the slopes of the lines joining

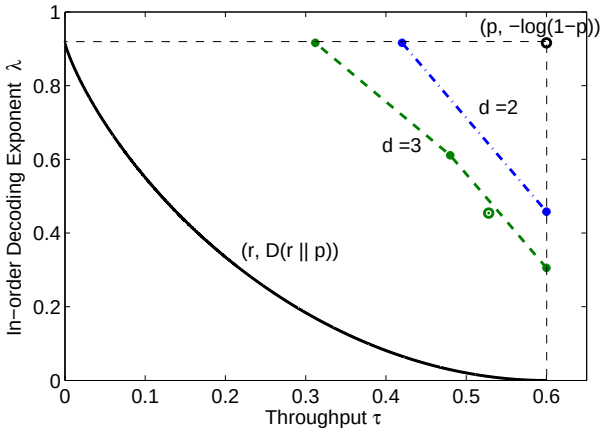

Fig. 4. The throughput-delay trade-off with $p = 0.6$ for $d = 2, 3$ which can be shown to be optimal over all convex combinations of time-invariant schemes. The point just below the piece-wise linear curve for $d = 3$, corresponding to the sub-optimal scheme $\mathbf{x} = [1, 2, 0]$.

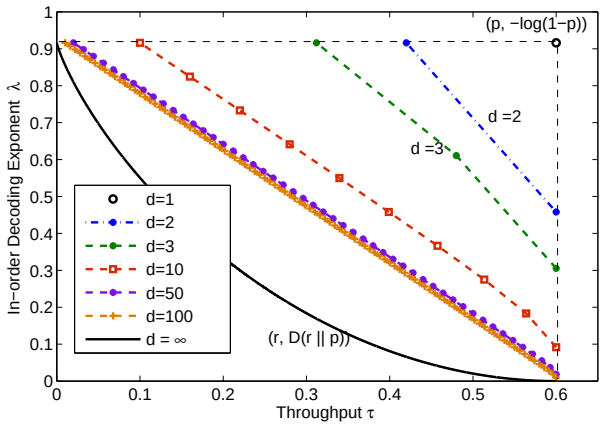

Fig. 5. The throughput-delay trade-off of the suggested coding schemes in Definition 5 with $p = 0.6$ and different values of feedback delay d. Numerical results suggest that this trade-off is optimal all convex combinations of time-invariant schemes for small d.

$(\tau_{\mathbf{x}^{(i)}}, \lambda_{\mathbf{x}^{(i)}})$ for $i = 1, \cdots 4$ satisfy,

$$\frac{\lambda_{\mathbf{x}^{(1)}} - \lambda_{\mathbf{x}^{(2)}}}{\tau_{\mathbf{x}^{(1)}} - \tau_{\mathbf{x}^{(2)}}} \geq \frac{\lambda_{\mathbf{x}^{(1)}} - \lambda_{\mathbf{x}^{(3)}}}{\tau_{\mathbf{x}^{(1)}} - \tau_{\mathbf{x}^{(3)}}} \tag{19}$$

$$\frac{\lambda_{\mathbf{x}^{(2)}} - \lambda_{\mathbf{x}^{(4)}}}{\tau_{\mathbf{x}^{(2)}} - \tau_{\mathbf{x}^{(4)}}} \leq \frac{\lambda_{\mathbf{x}^{(3)}} - \lambda_{\mathbf{x}^{(4)}}}{\tau_{\mathbf{x}^{(3)}} - \tau_{\mathbf{x}^{(4)}}}. \tag{20}$$

■

The trade-off for $d = 2$ and $d = 3$ with $p = 0.6$ is shown in Fig. 4. The point below the piece-wise linear curve for $d = 3$, corresponding to the sub-optimal scheme $\mathbf{x}^{(3)} = [1, 2, 0]$. We observe that the optimal trade-off becomes significantly worse are d increases. From this we can imply that frequent feedback to the source is important in delay-sensitive applications to ensure fast in-order decoding of packets.

For general d, it is hard to search for the $(\tau_{\mathbf{x}}, \lambda_{\mathbf{x}})$ that lie on the optimal trade-off. We suggest a set of time-invariant schemes which are easy to analyze and they give a good (τ, λ) trade-off.

Definition 5 (Suggested Schemes for General d)**.** *For general d we suggest schemes with $x_1 = a$ and $x_{d-a+1} = d - a$, for $a = 1, \cdots d$. They give the throughput-delay trade-off*

$$(\tau, \lambda) = \left(\frac{1 - (1-p)^a + (d-a)p}{d}, -\frac{a}{d} \log(1-p) \right). \tag{21}$$

Fig. 5 shows the trade-off given by (21) for different values of d. Observe that for $d = 2$ and $d = 3$ the suggested schemes coincide with the optimal trade-off we derived in Lemma 4 and Lemma 5 and shown in Fig. 4. As $d \to \infty$, and $a = \alpha d$, the trade-off converges to $((1-\alpha)p, -\alpha \log(1-p))$ for $0 \leq \alpha \leq 1$, which is the line joining $(0, -\log(1-p))$ and $(p, 0)$. Numerical results suggest that for small d this class of schemes gives the best trade-off among all possible time-invariant schemes $\mathbf{x}$, and close to optimal in general.

VI. Concluding Remarks

In this paper we analyze how block-wise feedback affects the trade-off between throughput τ and in-order decoding exponent λ, which measures the burstiness in-order packet decoding in streaming communication. When there is immediate feedback, we can simultaneously achieve the optimal τ and λ. But as the block size increases, and the frequency of feedback reduces, we have to compromise on at least one of these metrics. Our analysis gives us the insight that frequent feedback is crucial to ensure in-order packet delivery in delay-sensitive applications.

Given that feedback comes in blocks of d slots, we present a spectrum of coding schemes that span different points on the (τ, λ) trade-off as shown in Fig. 5. Depending upon the delay-sensitivity and bandwidth limitations of the applications, these codes provide the flexibility to choose a suitable operating point on trade-off. The proposed codes can be shown to be optimal over the broad class of full-rank codes for small feedback delay d, and when there is no feedback.

References

[1] Sandvine Intelligent Networks, "Global Internet Phenomena Report." http://www.sandvine.com, Mar. 2013.

[2] M. Luby, M. Mitzenmacher, A. Shokrollahi, D. Spielman, and V. Stemann, "Practical loss-resilient codes," in *ACM symposium on Theory of computing*, (New York, NY, USA), pp. 150–159, ACM, 1997.

[3] E. Martinian, *Dynamic Information and Constraints in Source and Channel Coding*. PhD thesis, MIT, Cambridge , USA, Sept. 2004.

[4] A. Badr, A. Khisti, W. Tan and J. Apostoupoulos, "Robust Streaming Erasure Codes based on Deterministic Channel Approximations," *International Symposium on Information Theory*, July 2013.

[5] P. Patil, A. Badr, A. Khisti and W. Tan, "Delay-Optimal Streaming Codes under Source-Channel Rate Mismatch," *Asilomar*, Nov. 2013.

[6] J. Sundararajan, D. Shah and M. Médard, "ARQ for Network Coding," in *International Symp. on Information Theory*, pp. 1651–1655, July 2008.

[7] J. Barros, R. Costa, D. Munaretto, and J. Widmer, "Effective Delay Control in Online Network Coding," in *International Conference on Computer Communications*, pp. 208–216, Apr. 2009.

[8] H. Yao, Y. Kochman and G. Wornell, "A Multi-Burst Transmission Strategy for Streaming over Blockage Channels with Long Feedback Delay," *IEEE Journal on Selected Areas in Communications*, Dec. 2011.

[9] G. Joshi, Y. Kochman, G. Wornell, "On Playback Delay in Streaming Communication," *International Symp. on Information Theory*, July 2012.

[10] T. Cover and J. Thomas, *Elements of information theory*. New York, NY, USA: Wiley-Interscience, 2nd ed., 1991.

[11] R. Durrett, *Probability: Theory and Examples*. Cambridge University Press, 4th ed., 2010.